AD-A007 560

REASONING FROM INCOMPLETE KNOWLEDGE

Allan Collins, et al
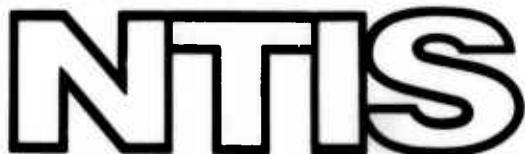
Bolt Beranek and Newman, Incorporated

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

March 1975

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report No. 4	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD-A007560
4. TITLE (and Subtitle) REASONING FROM INCOMPLETE KNOWLEDGE		5. TYPE OF REPORT & PERIOD COVERED Semi-Annual (1 Aug 74 - 31 Jan 75)
		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 3019
7. AUTHOR(S) Allan Collins, Eleanor H. Warnock, Nelleke Aiello, Mark L. Miller		8. CONTRACT OR GRANT NUMBER(S) No. N00014-71-0228
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton St., Cambridge, Mass.		10. PROGRAM ELEMENT PROJECT TASK AREA & WORK UNIT NUMBERS 61153N RR042-06;PRO42-06-01 NR154-330
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217		12. REPORT DATE March 1975
		13. NUMBER OF PAGES 50
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Education, CAI, Tutoring, Inference, Cognitive Psychology, Analogy, Induction, Causality.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The paper describes how people use a variety of plausible, but uncertain inferences to answer questions, about which their knowledge is incomplete. This kind of reasoning is described in terms of how it is being implemented in the SCHOLAR/CAI system. The paper also shows how people can be taught to reason in this way, using a Socratic tutorial method implemented in a system like SCHOLAR.

DD FORM 1 JAN 73 1473   EDITION OF 1 NOV 65 IS OBSOLETE

REASONING FROM INCOMPLETE KNOWLEDGE

Allan Collins

Eleanor H. Warnock

Nelleke Aiello

Mark L. Miller

Bolt Beranek and Newman Inc.

Cambridge, Massachusetts   02138

REASONING FROM INCOMPLETE KNOWLEDGE

Allan Collins
Eleanor H. Warnock
Nelleke Aiello
Mark L. Miller

Bolt Beranek and Newman Inc.

TABLE OF CONTENTS

INTRODUCTION

It doesn't trouble people much that their heads are full of incomplete, inconsistent, and uncertain information. With little trepidation they go about drawing rather doubtful conclusions from their tangled mass of knowledge, for the most part unaware of the tenuousness of their reasoning.   But the very tenuousness of the enterprise is bound up with the power it gives people to deal with a language and a world full of ambiguity and uncertainty.

We will describe this kind of human reasoning in terms of how a computer can be made to reason in the same illogical way.   For this purpose we will use SCHOLAR (Carbonell,   1970a,   1970b),   a computer program whose knowledge about the world is stored in a semantic network structured like human memory (Collins and Quillian, 1972). One of SCHOLAR's data bases is about geography, and people's knowledge about geography has the nice property, for our purposes, of being incomplete, inconsistent, and uncertain. So the examples and analysis will concern geography, but geography is only meant as a stand-in for everyman's knowledge about the world.

SCHOLAR's aim in life is to teach people by carrying on a tutorial dialogue with them (see Collins, Warnock, and Passafiume, 1975). Once upon a time, Socrates thought he could teach people to reason by such a tutorial method. We

will attempt to show that a person can learn to infer at
least some of what he doesn t know about geography by the
Socratic method, and to show how a program like SCHOLAR
might even play the role of Socrates with some finesse.

OPEN VERSUS CLOSED WORLDS

Recently Carbonell and Collins (1973) have stressed the
distinction between open worlds, such as geography, where
knowledge is incomplete, and closed worlds, such as the
blocks world of Winograd (1972) or the lunar rocks catalogue
of Woods, Kaplan, and Nash-Webber (1972), where the complete
set of objects and relations is known. The distinction is
important, because many of the procedures and rules of
inference that have been developed for dealing with closed
worlds do not apply to open worlds.

The distinction between open and closed worlds comes up
in a variety of ways. For example, if there are no basaltic
rocks stored in a closed data base, then it makes sense to
say "no" to the question "Were any basaltic rocks brought
back?". But, if no volcanoes are stored in a data base for
the U.S., it does not follow that the question "Are there
any volcanoes in the U.S.?" should necessarily be answered
"no". A more appropriate answer might be "I don´t know".
Furthermore, it makes sense to ask what the smallest block

in a scene is, but it makes little sense to ask what is the shortest river or the least famous lawyer in the U.S. It would be an appropriate strategy for deciding how many blocks in a scene are red, to consider each block and count how many are red. But it would not be an appropriate strategy to consider each person stored in a limited data base (such as humans have), in order to answer the question "How many people in the U.S. are over 30 years old?".

Within open worlds there are closed sets, however. For example, it may be possible to say how many states are on the Pacific, if they are all stored. Since closed sets are rare, it makes sense to mark the closed sets in memory rather than the open-ended sets. Then it is possible to apply closed-set strategies where the entire set is known.

The reason most sets are open is that most concepts are ill-defined. One could plausibly argue that there is a smallest city in the U.S., if we agree on some arbitrary definition of a city (e.g., incorporation by a state). But to use Wittgenstein's (1953) example, there is no way to specify precisely what is and is not a game. Even if we were to agree on some definition, we would get into difficulty when we try to apply it to cases. Outside of mathematics and logic, most concepts are simply not susceptible to precise definition.

Where a concept is relatively well-defined, like states in the U.S., we still may not know all the examples, and so we have to treat it as an open set.   This means that the distinction between open and closed sets is not in the outside world, but rather in each person's head.   Your closed set may be my open one.

We can illustrate some of the issues by considering Moldavia, since hardly anybody ever considers it, except perhaps Moldavians.   Most adult Americans know all the states in the U.S., so they know that Moldavia is not a state.   They may not be able to name all the states, but they've heard the states enough times that they've stored each of them as recognizably a state.   They may even know either explicitly (to name) or implicitly (to recognize) all the countries in South America well enough to say Moldavia is not one of those.

The same distinction between explicit and implicit knowledge exists in SCHOLAR.   The states would be stored implicitly if each appeared as an entry in the data base with an instance-of (superordinate) link to state. They would be stored explicitly if they were all stored as instances under U.S. states.

The same objects can be part of a closed set on some occasions and an open set on others. Even though a person (or SCHOLAR) may know all the countries in South America, he

may   not know all the countries in the world. So he may not
be able to say whether Moldavia is a country or not, even
though he can say it is not a country in South America.
Similarly none of us really knows whether Moldavia is a city
or town in the U.S., unless of course it is one. But by
restricting the set to, say, the major cities in the U.S.,
we can exclude Moldavia. Whether Des Moines is a major city
in the U.S. may be debatable, but there is no way Moldavia
can make it. Words like "major" or "typical" (Lakoff, 1972)
make it possible to restrict a set to exclude borderline
cases, such as the likes of Moldavia.

What it takes for a computer system like SCHOLAR to
discriminate between Moldavia and Des Moines are tags that
indicate the relative importance of different cities
(Carbonell and Collins, 1973; Collins, Warnock, and
Passafiume, 1975). Suppose there is a particular data base
configuration where a number of U.S. cities are stored,
with Moldavia not one of them and Des Moines tagged to be of
minor importance.   The decision rule as to what are the
major cities would be something like this: include those
that are tagged as important, exclude any not stored, and
any objects stored, but not clearly important, are excluded
or hedged about, depending on their relative importance and
the size of the set stored. In this way people can apply a
modified closed-set strategy to deal with open sets.

This strategy is just one rabbit from a seemingly open-ended hat. People have more such tricks than we can see, much less understand. There are negative tricks, functional tricks, visual imagery tricks, inductive tricks, and undoubtedly many more that people use to circumvent the holes and uncertainties in their knowledge. These all lie outside the deductive logic of which the advocates of theorem-proving and the predicate calculus are so fond.

NEGATIVE INFERENCES

People do not store most things that are not true, for example that Mexico has no king. Therefore, deciding that something is not true normally requires an inference. In a closed world, one can relegate whatever is not stored or not deducible from what is stored, to the dustbin of untruths. But in an open world, if one says "no" on that basis, then one will simply often be wrong. Therefore people use a variety of strategies to decide when to say "no", "probably not", "not really", or "I don't know".

Many of the strategies that people use to reach negative conclusions involve their functional knowledge, which we will discuss in the next section, or their visual knowledge, which we will not discuss in this paper. But there are several strategies we have been modelling in

SCHOLAR that are almost purely negative inferences.

Match Processes

Underlying the strategies people use to deal with their incomplete knowledge is a basic matching process for deciding whether any two concepts are the same or different. As we have argued elsewhere (Collins & Quillian, 1972), whether two concepts are treated as the same or different depends on the outcome of the matching process and the context in which it is called. The matching process pervades all language processing (Collins & Quillian, 1972) and we will refer to it repeatedly in the various inferential strategies described.

SCHOLAR does not mimic the human matching process perfectly (as described in Collins & Loftus, 1975) for reasons of computational efficiency, but it does capture the major aspects of that process. In SCHOLAR basic comparisons are handled by two processes called Match-on-Superordinate and Match-on-Superpart. Match-on-Superordinate decides "Can X be a Y?" and Match-on-Superpart decides "Can X be part of Y", for any two elements X and Y in the data base.

Both processes are basically intersection searches (Quillian, 1968, Collins & Quillian, 1972). They trace along the superordinate or the superpart links from both X and Y until an intersection is found. When an intersection

is found, they try to find a basis for saying "yes" or "no",
and only if they fail do they conclude "I don't know". The
flow   chart   in   Figure   1   shows   the   basic   strategy   of
Match-on-Superordinate   which   is   the   simpler   of   the two
processes.

We will explain how Match-on-Superordinate   works   in
terms   of   its   various   outcomes   in   answering   true-false
questions.   Examples 1 and 2 show different kinds   of   "yes"
outcomes.   (Expressions   between   asterisks   are   inputs to
SCHOLAR.) The first example involves a simple superordinate
inference,   and   the   second   shows   the   outcome   when   the
superordinate link is in the opposite direction   from   that
which the question presupposes.

(1)   *IS SANTIAGO A PLACE?*

      YES.   MORE SPECIFICALLY, SANTIAGO IS A CITY.

(2)   *IS A CITY A CAPITAL?*

      SOMETIMES.   MORE GENERALLY, A CITY IS A PLACE.

If there is no common superordinate of X and Y, a
reasonable   response   is   "no".   This   may   be the strategy
people use to say "no" when   two   things   are   semantically
unrelated.   For example, when asked if a cafeteria is a dog,
the reason they give for saying "no" is something   like   "A
cafeteria is one sort of thing and a dog is another sort of
thing".   In example 3 the   top-level   superordinate   for
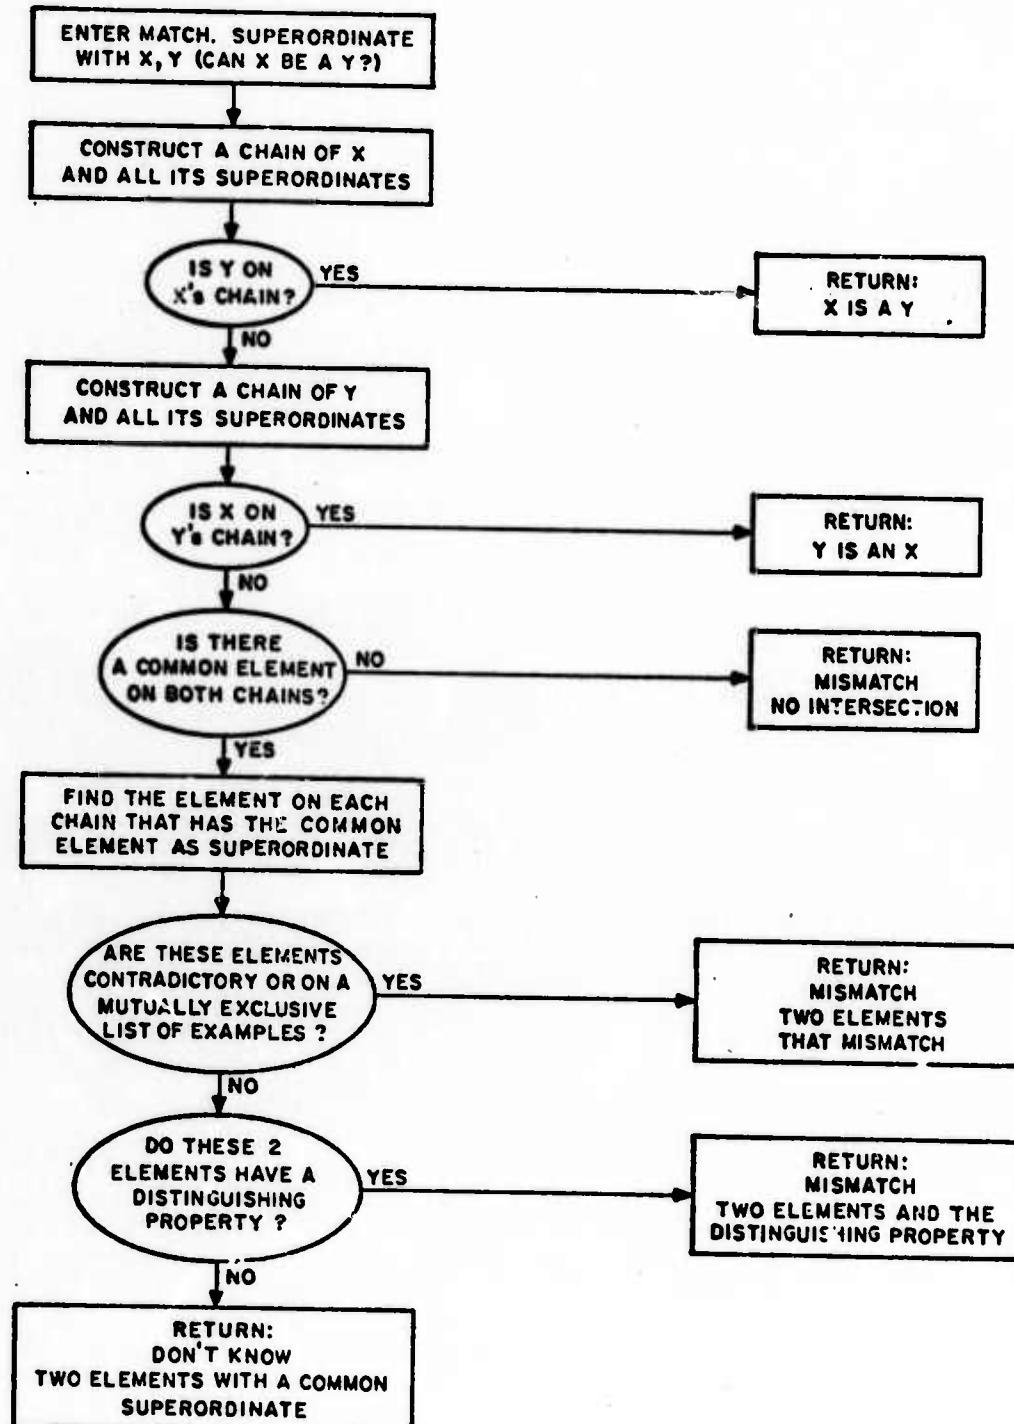
Figure 1.   Flow Chart for Match-on Superordinate

Surinam is place, which does not intersect with the chain for great man, and so the answer given is "no".

(3)   *IS SURINAM A GREAT MAN?*

NO.   SURINAM IS A COUNTRY, NOT A GREAT MAN.

If there is an intersection found, and the two things that led to the intersection are marked as mutually exclusive, then Match-on-Superordinate concludes "no". People may use just such a strategy to say "no" when asked if a sparrow is a wren. They probably don´t know any properties that are different between the two; they are just different kinds of birds. In example 4 Bolivia is a country and a capital is a city. City and country are mutually exclusive kinds of places, so Match-on-Superordinate concludes "no".

(4)   *IS BOLIVIA A CAPITAL?*

NO.   BOLIVIA IS A COUNTRY, NOT A CAPITAL.

The next example illustrates the case where the two objects, in this case Recife and Montevideo, have a common superordinate, but are not mutually exclusive. People frequently give a distinguishing property, such as a difference in location, as a reason for saying that two

things are not the same. This observation led to the test for a distinguishing property in Match-on-Superordinate. In example 5 Recife and Montevideo have the distinguishing property that they are located in different places, so the answer is "no". This difference is found by Match-on-Superpart.

(5)   *IS IT TRUE THAT RECIFE IS THE CAPITAL OF URUGUAY?*

NO.   RECIFE IS IN BRAZIL, NOT IN URUGUAY. THE CAPITAL OF URUGUAY IS MONTEVIDEO.

Example 6 shows the failure to find any basis for a distinction. A port can be a city and Sao Paulo is a city, and port and Sao Paulo are not stored as mutually exclusive. Nor are there any distinguishing properties between them. So there is no contradiction. This leads to a "Don't know" outcome.

(6)   *IS SAO PAULO A PORT?*

I DON'T KNOW IF SAO PAULO IS A PORT, BUT IT IS A CITY.

When a "Don't know" outcome occurs, it is appropriate to try a probabilistic inference, such as a lack-of-knowledge inference (Carbonell and Collins, 1973) or a functional inference. For example, it is possible to infer that Sao Paulo is probably not a port on the "lack-of-knowledge" basis that "I know other less important

cities that are ports and so I would know about it, if it were true".   Alternatively, the functional basis might be used that since it is not known to be on any major rivers or bodies of water, it must not be a port. These kinds of inferences will be discussed in more detail later.

Contradictions and the Uniqueness Assumption

Contradictions   appear   to   be   logically   certain inferences, but people's contradictions turn out to be uncertain inferences, based on incomplete knowledge. We can illustrate the uncertainty of contradictions with examples from actual human dialogues. The following examples show the basic contradiction strategy people use.

(Q)   Is Philadelphia in New Jersey?

(R)   No. It's in Pennsylvania, but it's across the river from New Jersey.

(Q)   Is Portuguese the language of Mexico?

(R)   No. Spanish is the language of Mexico.

The contradiction strategy that emerges from these two examples (as well as others) depends on meeting four conditions. The conditions are specified in terms of what is found or not found in memory.   In order to reach a contradiction to a query of the form "Is X in relation R to Y?" the memory search must meet the following conditions:

1) for all U that are found such that U R Y, U must be
distinct from (i.e.   mismatch) X, 2) for all V that are
found such that X R V, V must be distinct from Y, 3) for all
S that are found such that X S Y, S must be distinct from R,
and 4) either the Us or Vs must be a complete set.   The
first three conditions can be satisfied by failure to find
anything in memory (or by finding only some of the things
there) but the completeness condition (4) cannot. These
conditions are not at all obvious, and we will try to
explain them in terms of one of the examples.

.

The way these conditions must have been satisfied in
the Philadelphia example was as follows: 1) either he didn't
consider any places in New Jersey, or any he found (for
example Newark or Camden) must have been distinct from
Philadelphia, 2) the place he found Philadelphia to be in
was Pennsylvania, and he must have found that to be distinct
from New Jersey, 3) the   relation   he   found   between
Philadelphia and New Jersey was "across the river from" and
he must have found that to be distinct from "in", and 4) he
assumed that Pennsylvania was the only place Philadelphia
was in (i.e. that it was the complete set of locations for
Philadelphia). Though the first of these conditions did not
show up in the response in any form, it still must have been
met.   For example, if the respondent had known of a place
called "East Philadelphia" in New Jersey, his memory search
would probably have found it in this context (see Collins &

Quillian, 1972) and he would have had to decide if it matched Philadelphia or not. In such a case he would probably have hedged his answer with "Well, there is an East Philadelphia in New Jersey."

The uncertainty in this kind of inference arises for two reasons. The most obvious reason is that the memory search hardly ever finds all the Us, Vs, and Ss that are relevant to the decision. We think that the search for Us, Vs and Ss goes on in parallel, and the longer the question is considered, the more of them are found. What triggers a response is finding enough information to satisfy all the conditions. Sometimes, as we will see in a later example, a relevant piece of information is found after an initial answer has been given, and the person ends up changing his answer.

The other reason for the uncertainty lies in the default assumption of uniqueness that the respondents made in both examples, in order to meet the completeness condition (4). The uniqueness assumption has the form: If only one thing is found, then assume that it is the complete set. In the examples the respondents assumed that there is only one location for Philadelphia, and that Mexico has only one language, because these were the only values they found in memory. In fact, the assumption was wrong about Philadelphia; there is a Philadelphia, Mississippi, of civil

rights fame, and there may well be a Philadelphia in New Jersey.

The completeness condition need not be satisfied for sets of both Us and Vs. If one knows all the locations of Philadelphia, and they are all distinct from New Jersey, then it follows that Philadelphia cannot be in New Jersey. There is no indication that the respondent assumed, in the example, that he knew all the places in New Jersey, nor does that assumption ever appear to be the case.

In the examples below we will show how this basic strategy has been embedded in SCHOLAR. The first two are like the Philadelphia example in that the contradiction is based on information about the first object in the question. The mismatch occurs because Sao Paulo has only one location and there is a mismatch between it and the suggested location. In the first example the mismatch as to the location of Sao Paulo occurs on the basis of the country it is in, and in the second example on the basis of the region it is in. While the location is unique, it can be characterized in different ways, and a contradiction with any of those characterizations is grounds for saying "no".

(1)   *IS SAO PAULO IN CHILE?*

    NO. SAO PAULO IS NOT IN CHILE, BUT IN BRAZIL.

(2)   *IS SAO PAULO IN THE AMAZON JUNGLE?*

NO.   SAO PAULO IS NOT IN THE AMAZON JUNGLE, BUT IN THE
BRAZILIAN HIGHLANDS.

In the third example Sao Paulo might have been rejected
on the grounds that it is not among the cities in Chile, if
that set were complete.   (Completeness can be satisfied
either by making the uniqueness assumption or by marking the
set as exhaustive.) But the uniqueness assumption does not
apply,   because   there   is   more   than one city stored.
Furthermore, the set of cities is not marked as   exhaustive.
Therefore,   the   completeness condition is not satisfied for
cities in Chile.   But Sao Paulo is located in Brazil, not in
Chile,   and   Brazil   is distinct from Chile.   That is to say
the uniqueness assumption is applied to the location of Sao
Paulo, just as in the first example.   Example 4 is a
variation on example 3.   It illustrates the fact that the
contradiction need not be based on a mismatch with respect
to location.   It can   also   be based on a mismatch with
respect   to   the   superordinate (i.e., cities are not
mountains).

(3)   WHAT ARE THE PRINCIPAL CITIES IN CHILE?

SAO PAULO

NO.   THE   CORRECT   ANSWER   IS   SANTIAGO,   VALPARAISO,
CONCEPCION, ANTOFAGASTA, ARICA, VINA DEL MAR, PUNTA
ARENAS, AND IQUIQUE.   SAO PAULO IS NOT IN CHILE BUT IN
BRAZIL.

(4)   *IS ACONCAGUA A CITY IN PERU?*

NO.   ACONCAGUA   IS   A   MOUNTAIN,   NOT   A   CITY.   THE
PRINCIPAL   CITIES   OF PERU ARE LIMA, CALLAO, AREQUIPA,
IQUITOS, TRUJILLO, CHICLAYO, AND CUZCO.

The next two examples show the two basic variants for
finding   a   contradiction with respect to information about
the   second   object   in   the   question.   In   example 5   the
strategy is like that of the person who said that Portuguese
was not the language of Mexico.   Because there is only one
language   stored   with Uruguay, SCHOLAR assumes that Spanish
is the only language.   In example 6 there are   two   northern
bordering   countries   stored   for   Argentina, but the set is
marked as exhaustive, so the module can use this as a   basis
for saying "no".

(5)   *IS PORTUGUESE THE LANGUAGE OF URUGUAY?*

NO.   THE LANGUAGE OF URUGUAY IS SPANISH.

(6)   *IS   URUGUAY   THE   NORTHERN   BORDERING   COUNTRY   OF
ARGENTINA?*

NO.   THE NORTHERN BORDERING COUNTRIES OF ARGENTINA ARE
PARAGUAY   AND   BOLIVIA.   URUGUAY   IS   ONE   OF   THE
NORTHEASTERN BORDERING COUNTRIES OF ARGENTINA.

In the first dialogue excerpt we showed, the respondent
pointed   out   that Philadelphia is across the river from New
Jersey in order to explain the correct relation between   the
two   objects   in the question.   Like the human, SCHOLAR will
point out any other relation   between   the   two   objects   it
finds.   In example 6 the   search   turned up the fact that

Uruguay is the northeastern bordering country of Argentina, so, like the human, it added that information as clarification.

Example 7 shows the generality of the strategy for finding contradictions. Tropical refers to climate but nothing about climate is stored with Montevideo. However, Montevideo is part of Uruguay, and Uruguay's climate is temperate and maritime. There is a mismatch between tropical and temperate found by Match-on-Superordinate, and this is the basis for saying "no". This kind of nesting of inferences (here a part inference and a contradiction) is described in Carbonell and Collins (1973).

(7)   *IS MONTEVIDEO TROPICAL?*

       NO.   MONTEVIDEO IS TEMPERATE.

We have argued above that people often use the uniqueness assumption as a default assumption to meet the completeness condition. This can be seen most clearly in the dialogue shown below. The example is striking because it shows first how the uniqueness assumption operates to produce a contradiction, and then how it is overridden by finding more information.

(Q)   Is Springfield in Kentucky?

(R)   No.

(Q)   Why do you say it´s not in Kentucky?

(R)   Because I know where Springfield is.   It´s in Massachusetts.

(Q)   OK.

(R)   There might be a Springfield in Kentucky.   But I´m not really sure which one you´re talking about.

(Q)   Why didn´t you bring that up when I asked you the question?

(R)   Because I just assumed you were talking about Springfield, Massachusetts.

At the beginning of the dialogue the respondent was willing to say that Springfield was not in Kentucky, because it was in Massachusetts. But then she must have thought of another Springfield. (It is not uncommon to see people change their answers as they find more information in memory.) When she realized there was more than one Springfield and she didn´t know all of them, she gave a "Don´t know" kind of response. The reason she assumed the questioner was talking about Springfield Massachusetts, we would argue, is because that was the only Springfield she had thought of at first.

To see the extreme case of the uniqueness assumption, we recommend talking to a two-year-old. One two-year-old of our acquaintance, named Elizabeth, has been heard to respond to the accusation that she was a tease with the assertion "No, I´m a girl." This was striking because she didn´t know

what   a   tease   was.   She knew she was a girl, and anything
else had to be wrong.   With age, people become less certain.
It's hard to imagine that a man who was called a misogynist
and who didn't know what a misogynist was, would respond
"No, I'm a man". It is absurd because adults have learned
the multiplicity of things anyone can be. We suspect that
people   become   less   certain (and grow out of being
"sophomoric") as they become more knowledgeable, because
their greater knowledge leads to the storing of multiple
values and prevents them from   using   the   uniqueness
assumption as a default assumption with the kind of abandon
we see in our two-year-old.

The multiplicity of Elizabeth brings up the distinction
between multiple values which are not equivalent, and sets
(or lists) which are made up of equivalent elements.
Instances such as Elizabeth or the Amazon only have one
identity and one location at a time, in accord with the
current physics of our world. But this identity or location
can be described in a variety of ways. A person can be a
two-year-old, a girl, and a tease; and the Amazon can be in
South America, in Brazil and Peru, and even in the jungle.
Though these multiple values look like sets, they behave
differently from sets in some ways.   One   important
distinction is that any one value will suffice in answer to
a question or in making an inference. Thus for the location
of the Amazon, it is appropriate to say simply that it is in

South America or, alternatively, in Brazil and Peru. A set, on the other hand, is treated as a single element and should not be split into pieces. In the Amazon example, Brazil and Peru form a set, and so it is misleading to say simply that the location of the Amazon is Peru, just as it is misleading to say a zebra is black. It is not so bad, though, to say simply that the Amazon is in Brazil, because most of it is. When one or a few values within a set are predominant in importance, then they are often referred to as if they formed the complete set.

For the purpose of finding a contradiction, it is necessary to find a comparable element among the multiple values. Thus, to decide if the Amazon is in the desert, it is appropriate to say "no" because it is in the jungle. On the other hand, the reason the Amazon is not in Argentina is because it is in Brazil and Peru. Failure to find a comparable element was the trap into which the uniqueness assumption led our two-year-old friend.

It would be possible to store explicitly the general knowledge that the uniqueness assumption makes implicit. For example, we might have stored as a fact about countries in general that they have only one capital and one language (unless otherwise indicated). The trouble with this approach is that, like knowledge of syntax, this kind of knowledge does not seem to be something that people usually

know   explicitly.   For   example,   it comes as a surprise to
discover that while countries have multiple products,   mines
usually   have   only one product.   It is a generalization one
has to make from all the mines one has   encountered   in   the
past.   Thus such a scheme would lead to the storing of what
appear to be little-known facts.   While people may sometimes
st re   such relationships explicitly, we would argue that in
general this is implicit knowledge that is built into   their
inferential processes.

Lack-of-Knowledge Inferences

When they cannot find a contradiction, people sometimes
fall   back   on   what   we   have   called   a   lack-of-knowledge
inference (Carbonell and Collins, 1973).   This strategy   can
be seen in the dialogue excerpt below.

- (Q)   Are there any other areas where oil is   found   other
  than Venezuela? (the context is South America)

- (R)   Not particularly.   There is some oil offshore there,
  but   in general oil comes from Venezuela.   Venezuela
  is the only one that is making any money in oil.

In order to produce such a response we would argue   the
respondent must have: (1) failed to find oil stored with any
other South American country, (2) made a default   assumption
that   most   countries do not produce oil, (3) had oil stored
as fairly important for Venezuela, and (4)   had   information
stored   about   all   the   other   countries,   that   was   less

important than the importance of oil for Venezuela. On this basis the respondent concluded that if the other South American countries produced oil, he would know about it. He hedged his response in two ways: by saying "not particularly" rather than "No" and by saying Venezuela was the only one making money on oil. The hedges occur because at best this is an uncertain inference. In fact he was wrong; other South American countries were making money from oil, but not large amounts like Venezuela.

We can show how this strategy has been embedded in SCHOLAR with several examples. The lack-of-knowledge inference is called when there is a failure to find a contradiction. In such a case SCHOLAR compares the object in question to any similar objects that have the property in question. Suppose the question is about rubber in Guyana, as in the first example. SCHOLAR compares what it knows about Guyana and its agricultural products, with what it knows about countries that produce rubber (i.e. Peru and Colombia). In particular, it compares the depth of its knowledge about agricultural products of Guyana to the depth at which rubber occurs for Peru and Colombia. The more SCHOLAR knows about Guyana, the more certain it is that it would know Guyana produced rubber if it were true. Similarly, the more important rubber is for those countries that produce rubber, the more certain it is that it would know if Guyana produced rubber. In this particular case it

derived a moderate degree of certainty from the comparison.

(1)   *IS RUBBER AN AGRICULTURAL PRODUCT OF GUYANA?*

RUBBER IS AN AGRICULTURAL PRODUCT OF PERU AND COLOMBIA.   I KNOW ENOUGH THAT I AM INCLINED TO BELIEVE THAT RUBBER IS NOT AN AGRICULTURAL PRODUCT OF GUYANA.   THE AGRICULTURAL PRODUCTS OF GUYANA ARE SUGAR AND LUMBER.

Example 2 illustrates what happens when SCHOLAR cannot derive enough certainty from the comparison to say even a doubtful "No". In this case there were no minerals stored with the Chaco, so the comparison was based on the depth of its general knowledge about the Chaco. Because this comparison is less specific, less certainty can be derived from it. In this case, the result is a "don't know" response, in which SCHOLAR points out the similar object it found where nitrates are a product.

(2)   *ARE NITRATES MINERAL PRODUCTS OF THE CHACO?*

NITRATES ARE MINERALS OF THE ATACAMA DESERT.   BUT I DON'T KNOW ENOUGH TO SAY THAT NITRATES ARE NOT MINERALS OF THE CHACO.

If there are no similar objects that have the property in question, then a "don't know" response is appropriate. In example 3 nothing about a humid climate is found for the other mountain ranges in South America, so SCHOLAR in its clumsy prose admits the extent of its ignorance.

(3)   *IS IT TRUE THAT THE CLIMATE OF THE ANDES IS HUMID?*

I DON´T KNOW WHETHER HUMID IS A CLIMATE OF ANY
MOUNTAIN RANGE IN SOUTH AMERICA.

The lack-of-knowledge inference involves reasoning from
cases, like an induction.   One kind of induction has the
following form: if several objects in a class have a given
property, then an object of the same class that one knows
less about probably also has that property.   In contrast,
the lack-of-knowledge inference has the following form: if
several objects in a class have a given property, then an
object of the same class that one knows as much or more
about probably does not have that property. Thus it turns
out that the lack-of-knowledge inference is a kind of
inverse of induction.

FUNCTIONAL INFERENCES

People can often figure out what they don´t know by
reasoning from their knowledge about what it depends on. In
geography people´s knowledge about what depends on what is
almost always incomplete.   They often do not know all the
functional determinants that influence a given variable such
as climate, agricultural products, or population density.
They are even less likely to know precisely how the
different functional determinants (or factors) affect the
variable, the values for all the functional determinants, or
how the determinants interact. But from the protocols we

have looked at, people appear to reason like engineers: They make rough calculations using various default assumptions such as linearity, independence of factors, and normal values for unspecified factors. Then they make adjustments afterwards for perceived variations from these assumptions.

In this section, we will use excerpts from dialogues to illustrate these aspects of functional reasoning, as well as some of the different strategies people use in functional reasoning. All the excerpts are verbatim except for the last, which is reconstructed from notes.

The first example illustrates the form of people's functional knowledge; in particular, the temperature function and two of its functional determinants, latitude and ocean temperature. Here ocean temperature is treated as causing an adjustment of the temperature determined by latitude. What emerges from this and other examples is that temperature is regarded as a linear function of latitude, with adjustments for other factors like altitude, ocean temperature, and tree cover. These modifying factors are assumed not to affect the calculation unless they have unusual values. A person will never estimate the temperature of a place if he knows nothing about the latitude. But he may make a rough calculation of the temperature if he knows the approximate latitude but not the other factors, by assuming normal values for the other

factors (given no information to the contrary). This is true even though the variations in altitude (0 to 5 miles) affect temperature roughly as much as do variations in latitude. It is just that there is a clear default value near 0 in the distribution for altitude and none for latitude.

1. (T) Is it very hot along the coast here? (Points to Peruvian Coast)

    (S) I don't remember.

    (T) No. It turns out there's a very cold current coming up along the coast; and it bumps against Peru, and tends to make the coastal area cooler, although it's near the equator.

This example also illustrates another aspect of the storage of functional relationships: the distinction between general knowledge and specific knowledge. The general knowledge about temperature involves how it depends on various factors like latitude, altitude, and ocean temperature. The specific knowledge is information the tutor has stored about the fact that coastal Peru is cooler than comparable regions and about the cooling influence of the particular ocean current. The general knowledge is about "temperature" and the specific knowledge is about "the temperature of coastal Peru". A data base must, therefore, be able to have functional knowledge stored in both places, with pointers between the two indicating that the specific knowledge is a known instantiation of the general rule.

The second example shows a student answering both a "why" question (Why do they grow rice in Louisiana?) and a "why not" question (Why not in Oregon and Washington?). In answer to the first, the student mentioned only one functional determinant, the need for water. In the dialogues, people typically give only one or two reasons in answer to a "why" question, except when they have thought about the functional determinants previously. The reasons given are the matches found between the values stored for the particular place (in this case Louisiana) and the values required for the particular variable (in this case rice).

2. (T) Where in North America do you think rice might be grown?

   (S) Louisiana.

   (T) Why there?

   (S) Places where there is a lot of water. I think rice requires the ability to selectively flood fields.

   (T) O.K. Do you think there's a lot of rice in, say, Washington and Oregon?

   (S) Aha, I don't think so.

   (T) Why?

   (S) There's a lot of water up there too, but there's two reasons. First the climate isn't conducive, and second I don't think the land is flat enough. You've got to have flat land so you can flood a lot of it, unless you terrace it.

In answering the "why not" question in Example 2, he mentioned three of the four determinants of rice growing.

(He omitted soil fertility here, though it came up later.) A "why not" question in effect asks for any mismatches between the values required by rice and the values stored for the place in question.   It is very unusual in a "why not" question to mention a functional determinant, such as rainfall, where the value stored for the place matches the value stored for rice. In this case it happened because water supply was primed (Collins and Loftus, 1975) by the previous discussion.   That is in fact why the tutor picked Oregon, as we will discuss in the next section.

A mismatch on one factor is reason enough for not growing a given product, like rice. On the other hand, it is necessary to have matches on all the relevant determinants for a yes answer.   A correct answer to the first question about Louisiana would have mentioned all four factors.

In the third example the same student named three of the four functional determinants to answer why they grow wheat in the Plains.   (The fourth, terrain, is not so critical with wheat, so it is not surprising it was omitted.) Both wheat and rice growing occur over a range of temperature, so they are both threshold functions of temperature. For places on earth, rice growing has only one bound.   There are places that are too cold, but none that are too hot. On the other hand, wheat growing has two

bounds,   though   the   student was only concerned with one in his response.   There are places where it is   too   warm   for wheat,   as   well   as   too   cold   for   wheat.   Agricultural products, and, as we   shall   see,   population   density,   are typically   treated   as   threshold   functions   on the various functional determinants.

3.   (T)   They grow some wheat out in   the   plains.   Do   you
          have any idea why?

     (S)   Boy, these are questions for a city boy, you   know.
           For wheat, what do you need? You need fertile soil,
           and you need adequate rains, but not as much as you
           need for rice. You don't need a tropical climate
           for wheat.   They grow wheat way up in Canada with a
           shorter   growing   season.   So you need fertile soil
           and some rain, and at least some   section   of   time
           where the   temperature   doesn't   go   too far below
           freezing.

     In his response he mentioned that wheat needed   fertile soil   and   adequate   rains,   but not as much as you need for rice.   In people's talk about such   threshold   functions   as soil   and   rainfall,   they   only   use   fuzzy   values such as fertile and adequate.   We think it is important to   be   able to   represent   varying degrees of precision from the kind of values that appear in conversation to precise   numbers,   and to process either type as points on a continuum with a range of tolerance against which all   matches   or   mismatches   are evaluated.

     The fourth and fifth examples show how people can   make calculations   about   a variable, if they know the functional

determinants.   In example 4 the strategy for deciding whether rice is grown in Florida is to match Florida against all four functional determinants.   He mentioned that it matched terrain, and he may have figured out that Florida would match on temperature. He voiced reservations about the match on water supply, so it was a doubtful match to his mind. If he had considered the requirement for fertile soil, he might have rejected Florida for this reason. It turns out that rice is in fact <u>not</u> grown in Florida.   The fifth example shows a variation on the same strategy, where the student made a successful prediction. The procedure is to pick those places with the best overall match on all the functional determinants. In this case he was quite right about the Nile delta, and though he was more vague about the tropics, he was right as far as he went. These two examples show that functional knowledge gives people real predictive power, even though it is fallible.

4.   (T)   Do you think they might grow rice in Florida?

   (S)   Yeah, I guess they could, if there were an adequate
       fresh water supply.   Certainly a nice, big, flat
       area.

5.   (T)   What kind of grains do you think they grow in
       Africa, and where, then? (Pause) Well, where would
       they grow rice if they grew it anywhere?

   (S)   If they grew it anywhere, I suppose they'd grow it
       in the Nile region, and they'd grow it in the
       tropics where there was an adequate terrain for it.

The sixth example shows a tutor making a functional analogy with respect to cattle raising. He thought of a region, western Texas, that matched the region in Argentina called the Chaco in terms of temperature, rainfall, and vegetation, the functional determinants of cattle raising. Since he knew that western Texas was cattle country he inferred that the Chaco might be as well. A negative functional analogy might have occurred if the student had asked whether the Chaco produced rubber. Since the Amazon jungle and Indonesia produce rubber, the tutor could have said "no" on the basis of the mismatch between the Chaco and those regions, with respect to temperature, rainfall, and vegetation.

6. (S)   Is the Chaco the cattle country? I know the cattle country is down there.

    (T)   I think it's more sheep country. It's like western Texas, so in some sense I guess it's cattle country.

The last example shows another variation on the functional analogy. The analogy is between New Haven, for which the requested value was known, and New York City. The functional dependence used is that the number of piano tuners depends on population size. Probably the respondent did not have this particular functional dependence stored, but generated it, because he knew that it is people who use pianos and because he could figure out the ratio of

population sizes for the two cities. This is a particularly good example of the assumption of linearity (that the number of piano tuners increases linearly with population size) and a correction afterward of 15% to 30% downward for some deviation from the assumptions made. The adjustment might either have corrected for a perceived non-linearity (that the number of piano tuners, like members of Congress, does not quite increase linearly with population size), or for a perceived difference between New Haven and New York on another functional determinant (e.g., New Haven may be more cultural on the average than New York). What should be emphasized is that either kind of correction is applied afterwards, and that the second kind entails an assumption of independence of the two factors, population size and culture.

7. (T) How many piano tuners do you think there are in New York City?

    (S) Well there are 3 or 4 in New Haven, which has about 300,000 people. That's about one per 100,000. New York has about 7 million people, so that would make 70. I'll say 50 or 60.

These examples illustrate some of the various ways that people gain real inferential power from their imprecise knowledge about what depends on what. The next section shows how this kind of knowledge can be acquired.

LEARNING TO REASON

In Table 1 we show segments from a dialogue on population density. The tutor was the first author, and this is one of several dialogues discussing functional interrelationships in geography. These dialogues had the character of an inquisition, complete with mental torture.

What is most apparent from the dialogues is that the students were learning a great deal. The dialogue in Table 1 shows the most sophisticated of the students, and the student's learning in this dialogue is particularly obvious. The similarity to a Socratic dialogue is striking. What the students were learning was not so much facts about geography, but rather how to induce what is relevant and predict what is likely. In other words, they were learning to think like geographers.

TABLE 1

SEGMENTS FROM A TUTOR-STUDENT DIALOGUE ABOUT POPULATION DENSITY

T First, I am going to talk about population density. Where are the large densities in North America?

S In North America I would suppose the Northeast Corridor, Washington to Boston, would be the most densely populated area overall.

T Now, why do you suppose that is?

S Well, most of the air traffic passes back and forth between those places I believe. That's where you hear most of the problems about transportation.

T   No. That's a true statement, and what I want to know is why.

S   You want to know the proximate causes of it?

T   Yes.   The causes of why it is a true statement.

S   Well, there are all those cities there, right?

T   OK, why are the cities there?

S   H'm. Well, you get to the question of why are cities located   in   certain   places.   Well,   I   guess   for geographical and strategic reasons. New York is there, because it has the greatest natural harbor in the world, I hear. Ah. It was the place where our country was settled first and a lot of the immigration came here and a lot of the people tended to gravitate to those places. And political reasons, I suppose. Washington, being the capital of the country, attracts a lot of bureaucrats and professional people.

T   OK. Where else is the density high?

S   Well, working up from Washington, there's Baltimore.

T   No, I mean what other areas. You named the Northeast.

S   Other places that are dense would be the Chicago area.

T   Why do you suppose that's a dense area?

S   That seems like almost a meaningless question.   Because there's lots of people there.

                    (section omitted)

T   Now, do you have any feeling for why regions in China are densely populated?

S   Well, the proximate cause I suppose is lack of adequate birth control, and the population explosion.

T   Why didn't that happen in Siberia?

S   Yeah, there's probably a pretty strong   interaction between the birth control practices which have only now become even possible and the climate and food supplies of an area. Political factors are in there too. I suppose it's possible there could be a population explosion in Siberia, but it would just take a hell of a long time for it to get there.   You don't really start to get a population explosion unless there's an already adequate

population that keeps on growing   inexorably. · Then   it
starts to get...

(section omitted)

T   Why do you suppose Java has high population and the other
    Indonesian islands have low population density?

(section omitted)

S   Well,   I   would   doubt   there   would   be   large   cultural
    differences   between   the   islands,   although I think some
    parts are predominantly Hindu but most I think is Moslem.
    Neither   of   those sects are particularly strong on birth
    control.   Climate   differences   aren't   so   severe.
    Political - I think the seat of government is on Java.

T   But why is the seat   of   government   there?   Because   the
    people are there, right?

S   Yeah, maybe so.   It doesn't make much sense to talk about
    the   availability of ports in an archipelago.   There must
    be thousands of them for the taking.   Let's see,   there's
    climate   and   ports   and politics and food supply.   Maybe
    the soil is different on Java than it is on some   of   the
    other islands.

T   Hm hm, that's possible.

S   Maybe there's a difference in the   political   history   of
    that   island   and   the   others.   There   might have been.
    Other islands could have been part of different political
    organizations.   I think they used to belong to the Dutch
    - most of them did.

T   So did Java.

S   Yeah, most of it did, but maybe there   was   a   famine   on
    some island that wiped out a proportion of the population
    a few generations back.   That's pretty   hypothetical.   I
    would   just   suppose it had something to do with politics
    and food supply.   Not too much difference in climate.

T   Yeah, I don't think the politics   matter   really.   Yeah,
    well, I might mention that Sumatra has a very mountainous
    terrain.

S   Oh, the terrain.   Yeah, and the other place would be much
    flatter and better for rice growing and stuff.   Yeah.

T   You mentioned soil and you were hitting at it then.

(dialogue continued)

What the Student Learns

As the student progresses in the dialogue of Table 1, he accumulated a whole set of factors that affect population density.   He learned, from dealing with a   range   of instances,   what   were   the   important determinants of population density. It is a process of inducing general knowledge from specific instances.

His early difficulty in answering the question about why density is high on the East Coast and his complaints about the meaninglessness of such questions, indicate that initially he had no general knowledge stored about the reasons for population density per se. He did have specific knowledge about the density in different places, and about some of the reasons for that density. For example, he knew that New York had a good harbor and its port facilities made it a center of population. He also knew that immigrants had poured into the East Coast, and often settled there, and that people are attracted to where the government is.   But these were facts about New York and Washington which happened to be relevant to population density.   It was knowledge stored with specific instances, not information stored with population density explicitly.

In the course of the dialogue, he derived the following

factors   in   addition to those mentioned for the East Coast:
foreign trade from the West Coast, birth control from China,
climate   and   food   supplies   from   the   difference   between
Siberia and China, soil and terrain (the latter was   brought
up   by   the tutor) from Java, industrialization from Europe,
minerals from South Africa, and seafood   from   West   Africa.
As   he   accumulated   these   factors   they   became explicitly
stored as functional determinants of population   density   in
general.

When he was confronted with the problem of why Java has
a   high   population   density   and the rest of the Indonesian
islands generally do   not,   he   started   going   through   the
reasons   he   had   accumulated   to   see   if   he   could find a
potential difference between Java and the other   islands   on
the functional determinants of population density.   What the
inductive process had achieved up to this point was   not   so
much that the number of facts stored had increased, but that
the information had become stored with the   general   concept
of   population density.   It was now <u>available</u> for processing
with respect to Java.   Because of this, the student in   fact
gained   real   inferential   power.   The answer about Java and
the prediction   shown   in   an   earlier   section   about   rice
growing   in   the Nile delta are only two examples of how the
accumulation of functional knowledge enables the student   to
reason in a generative way from incomplete knowledge.

Not only did the student accumulate reasons for population density that were already stored with specific instances, but he also generated some new reasons. For example, he may have known that China has a large population because of its lack of birth control, but he probably did not have stored the fact that climate and food supply were also reasons why China has a large population density. He brought these up when forced to compare China with Siberia and to say why one has a large population density and the other does not. This is another, separate aspect of induction.

This induction process involves finding what mismatching properties of China and Siberia can produce a difference in population density. Obviously the fact that Siberia is a region and China is a country will not account for the difference in population density. But the differences in climate and food-growing capability both can, so these are what the student mentioned. A little later, in discussing India, the student revealed the connection he found between climate and population density. His idea was that people will die of exposure if the climate is too cold. Other possible connections between climate and population density are that people are attracted to warmer climates (which is why Florida has a large population) and that climate affects food-growing. He probably did not find the former connection, but the latter connection was probably

the basis for his bringing up food-growing.

We would argue that the connection between climate and population density is the result of an intersection process, like the one hypothesized by Quillian (1968, 1969). In the student's memory there must have been several different pieces of information: the fact that Siberia has a very cold climate, the fact that China has a moderate climate, the fact that prolonged exposure to cold leads to death, and the fact that death lowers population. Starting at Siberia, China and population density, the search had to find these four facts, which when taken together lead to a difference in population density between Siberia and China. Tying these facts together creates a new piece of information.

There were also a number of other things the students learned during the dialogues which we might enumerate briefly.

(a) They learned about second-order effects. When the student in the dialogue shown added food supply to his list of things that affect population density, this made it possible to see that soil and terrain for Java, and the proximity of the ocean in West Africa, might affect food supply and thereby affect population density. Thus the induction that food supply is a factor permitted the further induction of these

second-order factors.

(b)   They learned about the multiplicity of reasons for any
given fact. As we said, this student accumulated many
reasons for population density. If he had been asked
later why there is a large population density on the
East Coast, we think he would have included such
variables   as   climate,   food   supply,   and
industrialization.   This   is   shown   most   clearly in
Examples   3   and   4   in   the   section   on   Functional
Inferences, where initially a student gave only one
reason   for   rice   growing in Louisiana, but later gave
three reasons for wheat growing on the plains.

(c)   They learned about feedback effects and interactions
between different factors.   This student pointed out
(though not in the excerpt shown) a feedback effect
that   occurs   with   respect   to   capitals.   A capital
usually is located where the people are, but the fact
that the capital is there tends thereafter to attract
people to the area.   Interaction between factors showed
up in many cases in the dialogues. One such case was
that ports are only important for trade if there is
something to ship, which ties this factor to food
supply and industrialization.

In summary, the students were learning by induction, and the dialogues showed two different aspects to this induction process.   1) The students were deriving new functional determinants by comparison of contrasting instances, and 2) they were accumulating general knowledge about functional determinants from the specific instances. In both cases, the process involves gathering specific pieces of knowledge scattered about in memory and storing them together in a new configuration where they are more available.   This pulling together of old knowledge into new structures requires new interrelationships to be specified. It is the fundamental way new knowledge is created.

The Socratic Method of Teaching Geography

In the dialogues the tutor was following a strategy to force the student to think like a geographer.   The agenda for the discussions was simply to discuss the functional determinants of geographical variables such as population density and agricultural products, for different places on the five major continents.   There was no fixed set of questions to be asked.   But there was an a priori determination to ask "where" questions, "why" questions, and "why not" questions.   The "where" questions elicit what is stored as specific knowledge about the variable in question, or force a predictive calculation where nothing is stored directly.   The "why" and "why not" questions elicit whatever

reasons are stored explicitly, or force inductions. When the student could not answer a "where" question, the tutor usually provided the answers himself, and then asked the corresponding "why" or "why not" question.

During the dialogues, the tutor often picked a new place to ask about, and there was one strategy that he used systematically for picking a new place. This strategy in its most general form showed up near the beginning of the dialogue on population density (but not in the fragment shown). The student kept mentioning ports as a reason for population density. So the tutor asked about Mexico because the population density occurs mainly away from the ocean and the ports. Then he picked Alaska because there are a lot of potential ports and very little population density. The two places were chosen to force the student to see that ports were neither necessary nor sufficient for population density. This strategy is in essence the "near miss" strategy that Winston (1973) found was necessary for a computer program to induce concepts from instances of those concepts.

The "near miss" strategy occurred throughout the dialogues. Other examples were the selection of Siberia after China in Table 1 and of Oregon and Washington in Example 2 in the section on Functional Inferences. In the latter case, the student said that they grow rice in

Louisiana because there was a lot of water there.   This was
an incomplete answer in that it omitted the warm climate,
flat terrain, and fertile soil which are required for rice
growing. So the tutor picked as a "near miss" a place which
had the factor mentioned, (i.e. a lot of water), but which
did not grow rice. This was to make the student see that a
lot of water was not enough. The tutor was precluded from
picking a place where rice was grown and there was little
water, because water is necessary for rice growing.

There were two other aspects of the tutor's strategy
for picking places that emerged in the dialogues,
particularly with the less sophisticated students.   These
are basic aspects of the strategy to force the student to
learn from cases:

(a) The tutor picked well-known places with extreme values
    on important functional determinants. For example, in
    one dialogue on population density, he asked why they
    have a low density in places like the Sahara, Tibet, and
    Alaska. These places were brought up to draw out from
    the student lack of water, mountainous terrain, and cold
    climate as factors causing low density.   This is an
    effective strategy because it allows the student to
    derive functional determinants himself by dealing with
    cases where the relevant determinant is the most obvious
    explanation.

(b) The tutor picked different places with the same value on the functional variable (e.g. different places with high population density), where the value occurred for some of the same reasons and some different reasons. (This strategy parallels Winston's generalization cases.) For example, with population density, cases like Tibet and Alaska both involve cold climate, in one case because of the mountains and in the other case because of the latitude. This strategy is effective for two reasons.   First, by repeating factors the tutor can see if the student can apply what he has learned about one place to another place.   Second, by illustrating the different combinations of factors that lead to the same conditions, the student is forced to derive the most general form of the functional dependencies involved.

The major difficulty for a computer program to tutor by this method is for it to understand the answers by the student.   But this is not an insoluble problem, because the program does not have to understand the student very well. The program only has to see if the student has included those factors that the program knows to be relevant for the place in question. Teachers can read answers to questions on tests written in handwriting that they could not read otherwise. This is because they have a strong expectation as to what the answer should say. Similarly, in analyzing answers, the program can use its knowledge about what are

functional determinants and what are possible values for any particular place and for any variable like population density or agricultural products. In this way, the program can build at least a partial understanding of what the student is saying or not saying, even when his answers are ungrammatical or incoherent. The beauty of the Socratic dialogue is that a partial understanding is all that is necessary to guide further questioning. It is not altogether inhuman to carry on a conversation when you don't completely understand what the other guy is saying.

CONCLUSION

What emerges from this view of human inferential processing is that people can often extract what they do not know explictly from some forms of implicit knowledge by plausible but uncertain inferences. Cutting across the variety of strategies we have described, there are common aspects, in particular match processing and the various default assumptions people make. We would argue that these are basic elements common to all human reasoning, and that they are overlaid with a variety of heuristic strategies people have learned in order to give reasonable answers in the face of their incomplete knowledge.

REFERENCES

Carbonell, J.R.   Mixed-initiative man-computer instructional dialogues.   Ph.D.   dissertation.   M.I.T., 1970a.   (Also BBN Report No. 1971.)

Carbonell, J.R.   AI in CAI: An artificial intelligence approach   to   computer-aided   instruction.   IEEE Transactions on Man-Machine Systems, 1970b, MMS-11, 190-202.

Carbonell, J.R.   & Collins, A.M.   Natural semantics in artificial   intelligence.   Proceedings   of   Third International   Joint   Conference   on   Artificial Intelligence, 1973, 344-351.   Reprinted in American Journal of Computational Linguistics, 1974, 1, Mfc. 3.

Collins, A.M. & Quillian, M.R.   How to make a language user.   In   E. Tulving   and   W. Donaldson   (Eds.), Organization of Memory.   New York: Academic Press, 1972.

Collins, A.M. & Loftus, E.F.   A spreading activation theory of semantic processing.   Psychological Review, 1975, 82, No. 5.

Collins, A.M., Warnock, E.H., & Passafiume, J.J.   Analysis
    and synthesis of tutorial dialogues. In G.H. Bower
    (Ed.), The Psychology of Learning and Motivation,
    Vol. 9. New York: Academic Press, 1975.

Lakoff, G. Hedges: A study in meaning criteria and the
    logic of fuzzy concepts.   Papers from the Eighth
    Regional   Meeting,   Chicago   Linguistics   Society.
    Chicago: University of Chicago Linguistics Department,
    1972.

Quillian, M.R. Semantic Memory.   In M.   Minskey (Ed.),
    Semantic   information   processing.   Cambridge,
    Massachesetts: M.I.T. Press, 1968.

Quillian, M.R.   The Teachable Language Comprehender: A
    simulation   program   and   theory of language.
    Communications of the ACM, 1969, 12, 459-476.

Winograd, T. Understanding Natural Language.   New   York:
    Academic Press, 1972.

Winston, P.H. Learning to identify toy block structures.
    In   R.L.   Solso (Ed.), Contemporary issues in cognitive
    psychology: The Loyola Symposium.   New   York: Halsted
    Press, 1973.

Wittgenstein L. Philosophical Investigations (Translated by G.E.M. Anscombe). Oxford: Blackwell, 1953.

Woods, W.A., Kaplan, R.M. & Nash-Webber, B. The lunar sciences natural language information system. Bolt Beranek and Newman Inc., BBN Report No. 2378, 1972.